

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 485 690 A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 91109850.7

(51) Int. Cl.⁵: G06F 15/16, G06F 15/76

(22) Date of filing: 15.06.91

A request for correction of pages 30 and 31 has been filed pursuant to Rule 88 EPC. A decision on the request will be taken during the proceedings before the Examining Division (Guidelines for Examination in the EPO, A-V, 2.2).

(30) Priority: 13.11.90 US 611594

(43) Date of publication of application:
20.05.92 Bulletin 92/21(84) Designated Contracting States:
AT CH DE DK ES FR GB IT LI NL SE(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(72) Inventor: Dieffenderfer, James Warren
396 Front Street
Owego, New York 13827(US)
Inventor: Kogge, Peter Mivchael
7 dorchester Drive
Endicott, New York 13760(US)
Inventor: Wilkinson, Paul Amba
544 Forest Hill Rd.
Apalachin, New York 13732(US)
Inventor: Schoonover, Nicholas Jerome
Box 18, Route 17C
Tioga Center, New York 13845(US)

(74) Representative: Schäfer, Wolfgang, Dipl.-Ing.
European Patent Attorney, IBM Deutschland GmbH, Schönaicher Strasse 220
W-7030 Böblingen(DE)

(54) Parallel associative processor system.

(57) Multiprocessor parallel computing systems and a byte serial SIMD processor parallel architecture is used for parallel array processing with a simplified architecture adaptable to chip implementation in an air cooled environment. The array provided is an N dimensional array of byte wide processing units each coupled with an adequate segment of byte wide memory and control logic. A partitionable section of the array containing several processing units are contained on a silicon chip arranged with "Picket"s, an element of the processing array preferably consisting of combined processing element with a local memory for processing bit parallel bytes of information in a clock cycle. A Picket Processor system (or Subsystem) comprises an array of pickets, a communication network, an I/O system, and a SIMD controller consisting of a microprocessor, a canned routine processor, and a microcontroller that runs the array. The Picket Architecture for SIMD includes set associative processing, parallel numerically intensive processing, with physical array processing similar to image processing. a military picket line analogy fits quite well. Pickets, having a bit parallel processing element, with local memory coupled to the processing element for the parallel processing of information in an associative way where each picket is adapted to perform one element of the associative process. We have provided a way for horizontal association with each picket. The memory of the picket units is arranged in an array. The array of pickets thus arranged comprises a set associative memory. The set associative parallel processing system on a single chip permits a smaller set of 'data' out of a larger set to be brought out of memory where an associative operation can be performed on it. This associative operation, typically an exact compare, is performed on the whole set of data in parallel, utilizing the Picket's memory and execution unit.

EP 0 485 690 A2




FIG.1
Prior Art

Field of the Invention

The present invention relates generally to multiprocessor parallel computing systems and particularly to a byte serial SIMD processor parallel architecture and the system and methods for parallel array processing with a simplified architecture adaptable to chip implementation in an air cooled environment.

Cross reference to Co-Pending Patent Applications

Related commonly owned Co-Pending Applications also assigned to the same assignee as the present application at the time of filing include:

U. S. Patent Application 07/250,595, filed September 27, 1988, now abandoned in favor of its continuation application Serial No. 07/519,332, filed May 4, 1990 of James L. Taylor entitled "SIMD Array Processor" (which was originally published as EPO Application Serial No. 88307855/88-A on May 3, 1989); and

U.S. Patent Application 07/193,990, filed May 13, 1988, of H. Li, entitled "Methods and Circuit for Implementing an Arbitrary Graph on a Polymorphic Mesh"; and

U.S. Patent Application 07/426,140, filed October 24, 1989, of R. Jaffe et al entitled "Two-Dimensional Input/Output Scheme for Massively Parallel SIMD Computers"; and

U.S. Patent Application 07/439,758, filed November 21, 1989, of W.C. Dietrich, Jr. et al entitled "Method and Apparatus for Performing Memory Protection Operations in a Parallel Processor System".

Background of the Inventions

R.A. Stokes et al, U.S. Patent 3,537,074, issued October 27, 1970 described a array computer with parallel processors and a single programmable control unit, a plurality of registers for storing counterpart vectors, mask registers, and means responsible to a sequence of one or more control unit instructions for concurrently operating on data in the vector registers, and during the 1970's parallel processors described by Stokes became known as SIMD (Single Instruction Multiple Data) machines. Such machines may also be described as consisting of a programmable control unit driving an array of n parallel processors, each processor having a memory portion, an arithmetic unit, program decode portion, and an input/output portion. These systems were large boxes which were sometimes associated with a host mainframe. An important difference between the SIMD and the more common processors was that, within such systems all SIMD processors may have a different set of data in the associated processor, but all processors are governed by a common controller. SIMD computers are also different from a more common Von Neumann processor in that each instruction operates on a data vector rather than on a single operand.

The most common type of multiprocessor systems are Multiple Instruction Multiple Data systems where each processor may execute a separate program operating on a separate data set. The processors in a MIMD system may perform separate tasks or they may each perform a different sub-task of a common main task.

As the thinking about SIMD parallel processors progressed it was thought, as illustrated by U.S. Patent 4,435,758 issued to Lorie et al on March 6, 1984, entitled "Method for Conditional Branch Execution in SIMD Vector Processors", that they may be appropriate if the tasks of the system to be performed are highly independent and contention-free, but if the tasks have a contention for resources, then a network of synchronous processors working in a SIMD mode may be indicated. Indeed, U.S. Patent 4,435,758 described the problem and described improvements they had to make upon the developments of the R.A. Stokes et al 4,101,960 patent of July 18, 1978 to provide for conditional branch execution.

It has become the norm to describe the most advanced SIMD machines as synchronous bit serial processors arranged in an NxN matrix, as illustrated by an exemplary multi-processor system described in the paper by H. Li and M. Maresca entitled "Polymorphic-Torus Network", Proceedings of the International Conference on Parallel Processing, pp 411-414. Matrix vector multiplication methods for such a massively parallel architecture as described by H. Li et al of NxN processors connected physically by a mesh topology and with a mesh network overlaid by another switching network for reconfiguration purposes was described in detail in IBM Technical Disclosure Bulletin Vol. 32, No. 3A August 1989 for improving the speed of execution in the multiplication of a sparse matrix to a vector.

There does exist a 1990 publication which shows that thought has also been given to making SIMD and SIMD/MIMD machines which operate with multiple rows of processors, where all the processors in a given row are executing exactly the same instruction, but each row is programmed differently. See the H. Li, B.Dom and R. Jaffe example in IBM Technical Disclosure Bulletin, Vol. 32, 8B, January 1990 wherein is described an architecture with this configuration which was called a Parallel Local Operator Engine (PLOE) for processing specific repetitive memory inspection functions.

This present application envisions the development of an air cooled system which can be fabricated with existing technology which can perform on a chip set with a newly established architecture the tasks which have been thought assignable to SIMD machines, as well as tasks which have thought to be assignable to Von Neumann machines.

Upon a hindsight review of background art, it will be recognized that others have described the use of a processor with a ROM for static instructions and registers for storing and accessing data on a single silicon chip, see for example the U.S. Patent No. 4,942,516 issued July 17, 1990 to G. Hyatt and entitled "Single Chip Integrated Circuit Computer Architecture", but this kind of work was unrelated to SIMD type complex applications.

Others have described various means for different tasks to which our system may also be directed. It is for instance known that matrix multiplication can be parallelized. There are publications which relate to applications to which our system may be employed which relate to Artificial Intelligence. Content addressable (Refer to FIGURE 3) or associative memories have been addressed at the university level with various processing chips. There it was learned that for AI applications in certain cases it is valuable to base row selection on the results of previous search operations, with row logic. See VLSI for Artificial Intelligence; Jose G. Delgado-Frias and Will R. Moore, editors, Kluwer Academic Publishers, 1989, pp. 95-108, "VLSI and Rule-Based Systems", Peter Kogge, Mark Brule, and Charles Stormon for a discussion of applications. However, the elegant suggestions of others seem to direct the art in a different direction than we have followed. One form is that of Oxford's "Intelligent-Memory" chips for massive parallelism which was described in VLSI Systems Design, December 1988, pp. 18-28 in an article by the Senior Editor, Bob Cushman entitled Matrix Crunching with Massive Parallelism. Others have attempted parallel processing on VLSI Associative memory in order to describe a VLSI associative memory chip suitable for the implementation of a reconfigurable fully parallel associative processing scheme. This approach thought that use of a classical associative memory architecture would require an inordinate number of pinouts for the data transfers involved. See e.g. Parallel Processing on VLSI Associative Memory, S. Hengen and I. Scherson, a report supported by an NSF award #ECS-h404627 and reported by the authors at the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106.

The problem which we have faced is the need to make compact processors for complex applications, and in approaching this general problem we have found that past designs which have been restricted to bit serial implementations with up to a few thousand bits of memory per processing element, and a few processing elements per electronic chip were inadequate. As a result we came upon the need to define an architecture which could be manufactured with dramatically higher densities and yet be suitable for an air cooled environment and could be packaged in a light and compact unit.

As a result, we came upon a solution which entails an N dimensional array of byte wide processing units each coupled with an adequate segment of byte wide memory and control logic. A partitionable section of the array containing several processing units are contained on a silicon chip. This will be described in detail below. Our own solutions generated additional problems to be solved to improve performance of the system and we have invented solutions to further the development of the propounded architecture. Some of these solutions enable independent performance of processing elements. In this connection under hindsight review we recognized that the prior art described in European Patent Application EP-A-208,457 described a processor array in which each processing element in the array was enabled to select the element from which it takes its input.

Co-Pending Application 07/519,332 of James L. Taylor, filed May 4, 1990, continuing U.S. Serial No. 07/250,595, filed September 27, 1988 and entitled "SIMD Array Processor" describes a multi-dimensional array of processing elements which has an enhanced degree of flexibility to enable the potential for parallel processing to be better exploited without resorting to the expense and complexity of a MIMD processor. This application was originally first published on May 3, 1989 in the form of an European application EPO #88/307885/88-A. The system there described sends a global instruction for local bit serial execution along a bus connecting control logic in the various parallel processing elements and programmatically modifies selected bits of the global instruction for use on local bit lines where the modified bits are decoded.

Hindsight review has also identified a single discussion of byte wide SIMD processor elements entitled "Design of SIMD Microprocessor Array" by C.R. Jesshope, R. O'Gorman, et al and published in IEEE Proceedings Vol. 136, May 1989. The paper describes a scholarly study of SIMD architectures by authors who are with the Department of Electronics and Computer Science at The University, Southampton, UK. Several features are similar to the matters we discuss, and they describe a processor which has a byte wide SIMD architecture. The paper suggests processing elements with an eight bit processing accumulator architecture elements, a limited on board RAM (1k bytes) and multiple PEs per chip (4) as well as some significant degree of local autonomy. However, in spite of the use of such a suggested structure, associative

processing is not provided. It will be seen that the structure suggested does not provide byte wide neighbor communication, as is the case of the present invention. Other features which we have developed are also not suggested. Accordingly, a need exists for our Picket system which is executable on an integrated circuit chip in order to advance the state of the art for parallel processing. See - A "Summary of Parallel Computer Architectures", R. Duncan, which may be found in the February 1990 IEEE Computer.

Definitions

For ease in following the description, here follows a general, non-limiting description of new terminology in the art.

Picket - an element of the processing array preferably consisting of combined processing element with a local memory for processing bit parallel bytes of information in a clock cycle. The preferred embodiment consisting of a byte wide data flow processor, 32k bytes or more of memory, primitive controls, ties to communications with other pickets.

Picket Chip - contains a plurality of pickets on a single silicon chip.

Picket Processor system (or Subsystem) - total system consisting of an array of pickets, a communication network, an I/O system, and a SIMD controller consisting of a microprocessor, a canned routine processor, and a microcontroller that runs the array.

Picket Architecture - the preferred embodiment for the SIMD architecture with features that accommodate several diverse kinds of problems including:

- set associative processing
- parallel numerically intensive processing
- physical array processing similar to images

Picket Array - a collection of pickets arranged in a geometric order.

The term "picket" comes from Tom Sawyer and his white fence, although it will also be understood functionally that a military picket line analogy fits quite well.

Summary of the Inventions

The architecture which we will elaborate may be referred to as a "picket" architecture. It incorporates certain processor element structures similar in some respect to the earlier SIMD processing elements (PEs) which march together when desired like a picket line, but which have been improved and are capable of independent Von Neumann operation. This architecture has been developed to employ an N-dimensional array of processing units arranged for parallel processing and having control logic within each processing unit which we call a picket.

The preferred system is capable of computing like a Single Instruction Multiple Data (SIMD) streaming machine, but performance has been improved by more parallel processing element implementations. The problems encountered with data dependency have been eliminated. In SIMD machine operations, no processor or function can have any data dependence in it that would cause one processing element to require a different number of cycles.

The architectural developments we have made resulted in a system with a plurality of picket units, generally Pickets, having a bit parallel processing element, with local memory coupled to the processing element for the parallel processing of information in an associative way where each picket is adapted to perform one element of the associative process. We have provided a way for horizontal association with each picket. The memory of the picket units is arranged in an array. The array of pickets thus arranged comprises a set associative memory.

The invention of a set associative parallel processing system on a single chip permits a smaller set of 'data' out of a larger set to be brought out of memory where an associative operation can be performed on it. This associative operation, which is typically an exact compare, is performed on the whole set of data in parallel, utilizing the Picket's memory and execution unit.

In the Picket array, each picket has a portion of data out of the larger set. Additionally, each picket selects one piece of data from it's portion. Thus, one piece of data in each of a set of pickets comprise the set of data upon which an associative operation is performed by all pickets in parallel.

The design which can be implemented today with up to 50k gates of data-flow and control logic, and with multi-megabits of DRAM memory on a single chip. Each of these chips is configured to contain plurality of pickets or processing units. In the preferred embodiment for text processing, which is capable of graphics use, there are 16 pickets with 32kbytes of DRAM memory for each picket on a single picket chip, and the system comprises an array of 64 of such picket chips, yielding an array of 1024 processing

elements. This picket architecture may be manufactured with CMOS technology which permits 4 million bits of DRAM be laid down in repetitive patterns on a single chip, and the remaining chip surface area available on the chip is filled with logic standard cells which will employ up to 50k logic elements and which can form the data flow and control logic in order to form pickets on the chip. We have arranged the system so that pickets may process data with local autonomy and provided a "slide" between pickets.

The picket technology is expandable, and with 128k byte of DRAM in each single picket (16 MBit DRAM memory chip), the picket architecture will handle full 24 bit color graphics in the same manner as text and 8 bit color or gray scale graphics are handled with our current preferred embodiment. Experimental manufacturing technology shows that this density is foreseeable within the near future as a consistent manufactured product capable of operating in an air cooled environment. For color graphics, our preferred picket architecture would increase the amount of DRAM on the chip to 128 kbyte per picket, while maintaining 16 pickets per chip. Alternatively, 24 picket units per picket chip with 96 kbyte memory could be employed for full color graphics processors.

15 Description of the Drawings

We will describe our preferred embodiment in relation to the accompanying drawings in which:

- Fig. 1 is a schematic illustration of a recent SIMD processor which is described in a copending application which may be considered as representative of the prior art.
- 20 Fig. 2 illustrates a pair of the basic picket units which are configured on a silicon base with processor, memory, control logic and associative memory with byte communication with the other pickets in the array.
- Fig. 3 illustrates associative memory processing.
- Fig. 4 illustrates a basic 16 (n) picket configuration for a SIMD subsystem which employs a microprocessor controller, a hardwired sequencing controller for canned routines, and a picket array and which forms the basic parallel picket processor system which may also be a stand alone unit.
- 25 Fig. 5 illustrates a multiple picket processor system which incorporates a plurality of picket processors of Fig. 4.
- 30 Fig. 6 is a sub-system functional block diagram.
- Fig. 7 further illustrates a subsystem controller's arrangement with the cards of Fig. 5.

Detailed Description

35 Turning now to the drawings in greater detail, it will be recognized that Fig. 1 represents a typical prior art SIMD system, of the type generally described in Taylor's IBM European Patent Application No. 88307855/88-A and in UK-A-1,445,714. In such prior art devices, the SIMD computer is a single instruction, multiple data computer having a parallel array processor comprising a plurality of parallel linked bit serial processors each being associated with one of a plurality of SIMD memory devices. The input/output I/O system acts as a staging system to the SIMD unit, and comprises a temporary store for the bi-directional 40 two-dimensional transfer of data between the host computer (which may be a mainframe or a microprocessor) and the SIMD computer. The I/O system includes input output processing means for controlling the flow of data between the host computer and the temporary storage means and for controlling the flow of data between the temporary store and a plurality of SIMD memory devices which are usually organized 45 buffer sections or partitions of a larger memory. Thus the input operation of the I/O system involves the transfer of data from the host computer memory to the temporary store and from the temporary store to the SIMD memory devices in the second step, and for output there is also a two step process which transfers across the two dimensional bus the data between the host computer and the SIMD computer. The input/output system for I/O transfer may be a separate unit, a sub-unit in the host, or often a unit within the 50 SIMD computer where the SIMD controller acts as the control for the temporary I/O buffer store.

The SIMD computer itself comprises a processor array having a plurality of processing elements and a network which connects the individual processing elements and a plurality of conventional separate SIMD memory devices. The SIMD computer is a parallel array processor having a great number of individual processing elements linked and operated in parallel. The SIMD computer includes a control unit that 55 generates the instruction stream for the processing elements and also provides the necessary timing signals for the computer. The network which interconnects the various processing elements includes some form of interconnection scheme for the individual processing elements and this interconnection can take on may topologies such as mesh, polymorphic-torus and hypercube. The plurality of memory devices are for the

immediate storage of bit data for the individual processing elements and there is a one-to-one correspondence between the number of processing elements and the number of memory devices which can be the aforementioned buffer partition of a larger memory.

For example, as illustrated in Fig. 1, there is provided a host processor 28. This processor is used to load microcode programs into the array controller 14 (which includes a temporary store buffer) to exchange data with it and to monitor its status via a host-controller data bus 30 and an address and control bus 31. The host processor in this example could be any suitable general purpose computer such as a mainframe or a personal computer. In this prior art example, the array of processors of the array are illustrated as on a 2-D basis, but the array could be organized differently, as on a 3-D or 4-D cluster arrangement. The SIMD array processor comprise an array 12 of processing elements $P(i,j)$, and an array controller 14 for issuing the stream of global instructions to the processing elements $P(i,j)$. While not shown in Fig. 1, the prior example has processing elements which operate on a single bit at one time and has associated there a block of storage which is a partition in memory associated with the processing element. The processing elements are connected by a so called NEWS (North, East, West, South) network to their respective neighbors by bidirectional bit lines. Thus, the processing elements $P(i,j)$ is connected to the processing elements $P(i-1,j)$, $P(i,j+1)$, $P(i,j-1)$, and $P(i+1,j)$ in the Northern, Eastern, Western and Southern directions respectively. In this typical example, the News network is toroidally connected at its edges so that the Northern and Southern Edges are bidirectionally interconnected and the Western and Eastern edges are similarly interconnected. In order that data may be input to and output from the array of processors, a controller-array data bus 26 is connected to the NEWS network. As shown it is connected to the East-West boundary of the array. It could be connected instead or additionally to the North-South boundary by means of bidirectional tristate drivers which are connected to the toroidal East-West News connection. As in the preferred embodiment which will be described, 1024 processing elements would be achieved by the prior art if the number of processing elements were in this example 32×32 instead of the illustrated 16×16 . In the illustration a single line indicates a single bit line, while a double line connecting functional elements is used to represent a plurality of connection lines or a bus.

In this prior example the array controller issues instructions in parallel to the processing elements via an instruction bus 18 and issues row select and column select signals via row select lines 20 and column select lines 22, respectively. These instructions cause the processing elements to load data from storage and to process the data and then to store the data once more in storage. For this purpose each processing element has access to a bit slice (section or buffer) of main memory. Logically therefore, the main memory of the array processor is separated into 1024 partition slices for a 1024 processing element array. This means that up to thirty-two 32 bit words can be transferred in or out of storage at one time in a transfer step. To perform a read or write operation, the memory is addressed in terms of an index address which is supplied to the memory address lines via an address bus 24 and read or write instruction is supplied to each of the processing elements in parallel. During a read operation, the row and column select signals on the row and column select lines identify which is the processing element that are to perform the operation. Thus in the described example, it is possible to read a single 32 bit word from memory into the thirty-two processing elements in a selected row when the array is 32×32 . The processing element is associated with the slice, or block of memory (i,j) which is one bit wide. While the slice or block memory is logically associated on a one-on-one basis with the associated individual processing element it may be and typically is physically separated on another chip. We do not know how with this prior architecture the described array processor could be manufactured, as can our picket be manufactured with an array of processors and adequate memory on a single chip of the kind we describe below.

Processing elements $P(i,j)$ of the prior example themselves should be understood to comprise an ALU, with input and output registers which will include a carry, each of which is able to store a single bit of information. There is a multiplexer which is connected to the ALU input, output and also to the bidirectional data port of the slice of memory (i,j) associated with the individual processing element $P(i,j)$.

There are separate instruction and data busses, and the array controller has a microcode store in which the microcode defining the processing to be performed by the array is loaded by the host 28 using the data bus 30 and the address and control bus 31. Once the operation of the array controller has been initiated by the host 28, the sequencing of the microcode is controlled by the microcode control unit which is connected to the microcode store within the array controller 14. An ALU and register bank of the array controller are used in the generation of array memory addresses, loop counting, jump address calculation and general purpose register operations which are output on an address bus of the array controller. The array controller also has mask registers for decoding row and column mask codes, and the specific operation codes pass to the processing elements over an instruction bus. In this example the array controller could have a data buffer within the controller but functionally between the host controller data bus and the controller array data

bus. From this buffer the data is loaded under control of the microcode in the control store into the array of processors, and vice versa. For this purpose the buffer is arranged as a bidirectional FIFO buffer under control of a microcode control in the array controller. More details of such prior systems may be found by reference to the above cited examples, particularly in Taylor's SIMD Array Processor, U.S.S.N. 07/519,332.

5 A review of the foregoing prior attempts may be compared to the preferred embodiments of the inventions described herein. Fig. 2 illustrates the basic picket unit 100, comprising a combined processing element ALU 101 with a local memory 102 coupled to the processing element for processing one byte of information in a clock cycle. As illustrated, the picket unit is formed on a silicon base chip or Picket Chip with a linear array of Pickets with neighbors to the side (to the left and right in the figure) such that on a
10 silicon base chip there is formed a picket processing array with a plurality of local memories, one for each of the byte wide processing data flows arranged in a logical row or in a linear array with neighbor communication busses for passing data bidirectionally to the right and left. The collection of pickets in a Picket Chip is arranged in a geometric order, preferably horizontally on the chip. Fig. 2 shows the typical implementation of two pickets of the picket array on a picket chip with multiple memory and dataflow
15 including communication paths between each picket's processing element and memory. In our preferred embodiment, the data communication paths between the one-on-one memory with the processing elements of the array is byte wide and across, left or right with a neighbor or with a "slide" for communication with picket processors farther away.

A "slide" may be defined as means for transferring information in a single cycle to a non-neighbor
20 position through a picket address location which would normally be able to receive the information were it not transparent to the message being sent until it arrives and is received at the nearest active neighbor which receives it. Thus a slide functions by sending information to a non-neighbor position across "turned off" pickets. Say picket "A" wants to transfer information to a remote picket "G". Prior to that cycle, intervening pickets are made transparent by turning these "B" through "F" pickets off. Then in the next
25 single cycle "A" sends his message to the right and in doing so passes through "B" through "F" which are transparent because they are turned off and "G" receives the message since it is still turned on. In the normal use of the "slide" information is transferred linearly across the lattice, but the slide approach can also work with a two dimensional mesh, or in a multi-dimensional array.

The access of the processing elements in our preferred embodiment is not bit serial in operation, but
30 rather byte serial. Each processor has access to its own coupled memory, rather than having access to a block of local memory and an associated partition or page thereof. Instead of one bit, a character wide, or character multiples wide bus is provided. Instead of a bit, a byte (or in future systems envisioned to duplicate the performance of a character byte, multiple bytes) of information is processed in one clock cycle. Thus, 8, 16 or 32 bits may flow between each picket processing element to match the width of the
35 associated memory. In our preferred embodiment each picket chip has 8 (9) bits wide memory of 32 Kbytes and preferably 16 pickets with this 32 Kbytes each of storage per picket node of the linear array. In our preferred embodiment, each associated memory is founded in CMOS as DRAM, and a character byte is 9 bits (which functions as an 8 bit character with self checking).

The parallel path byte wide bus data flow between pickets and between the processing element and its
40 memory is a substantial improvement over the serial bit structure of the prior art systems, but it will also be recognized after this achievement is recognized that the increased parallelism results in additional problems which need solution as one grows into the understanding of the implications of the newly achieved architecture. Important solutions are described herein.

A feature which will be appreciated is that in addition to the left and right neighbor transfers, and the
45 slide mechanism which we have described with reference to the drawings, we have provided also a broadcast bus which is double byte wide, so that all pickets can see the same data at the same time. Picket control and address propagation is also transferred on this broadcast bus. It is this bus that supplies the comparison data when performing set association operations, and other comparison or synchronous math operations.

50 The tasks that have highly parallel data structures that lend themselves to processing within the picket data processing elements under the control of a single instruction stream include applications in artificial intelligence pattern matching, sensor and track fusion in multi-sensor optimal assignment, context searching and image processing. However, many of these applications now possible were not used in SIMD processes because the serial bit processing under a single clock time. For example, the traditional serial
55 processing element of a SIMD machine executes one bit of an ADD operation for each processor cycle, while a 32 bit parallel machine can execute 32 bits of an ADD in one cycle.

The configuration of 32 KB per processing element applies far more memory logically available to each processing element than the traditional SIMD machine provides.

Pin count on the chip has been held low because the data which is passed into and out of the chip has been kept to a minimum. The DRAM memory is a conventional memory CMOS array and one that supports "row-column" access by deleting the column demultiplexing on the back of the memory array, and providing a row address that reads out a row of the memory array to the data flows in parallel.

5 The memory, in addition to data, contains "tri-bits" or "trit", so that there are three states recognized by the logic, instead of the traditional binary digit, either logic 1, logic 0, or don't care. The don't care in a match field matches either a logic 1 or logic 0. The trit is contained in successive storage locations in the storage array. Masks are another form of data stored in memory that is directed to the mask register of the picket processing element.

10 As the storage array can contain commands, this allows one picket to do a different operation from another picket. On-chip control of individual pickets during operations involving most pickets, but not necessarily all, allows implementations which are unique for SIMD operations. One simple control function provided is that of suspended operations in any picket whose status output meets a specific condition. Thus a non-zero condition may mean a doze. Doze is a condition which suspends operations and turns the picket into an inactive but aware status. Another command provided is to inhibit or enable write to memory based on conditions in the picket, or based upon a command provided to the bus prior to a slide operation.

By applying to a picket chip 16 powerful pickets each with 32 KB memory, only 64 chips provides 1024 processors and memory of 32768 KB. The array of pickets comprises a set associative memory. The inventions are also useful for numerically intensive processing image analysis as well as vector processing. 20 This powerful picket processing array can be packaged today on only 2 small cards! It will be seen that thousands of pickets can be appropriately packaged in a more portable low power package, enabling image processing applications to be performed with minimal delay or within video frame time, for instance, during the flight of an airborne craft without significant payload consideration.

The power of the pickets leads to the possibility of use of large associated memory systems packed 25 into close quarters and enables the use of the processing power in a variety of applications after the systems designer becomes accustomed to use of the new system.

Fig. 3 illustrates what might be termed fully associative memory in that, when an association is requested, a compare value is presented to all memory locations and all memory locations simultaneously respond with their match lines. Associative memory is known by itself in the art. In the system herein 30 described, using parallel pickets of memory and processing elements which have byte transfers in order to perform a search, there is an input of data and a mask for a search in order to locate a word K among N words in memory. All matching pickets raise a status line, and then a separate operation reads or selects the first match K. This operation, commonly called set associative, can be repeated for successive words up thru the picket memory. Similarly, writing is achieved via a broadcast operation in which a raised select line 35 indicates participation and broadcast data is copied to all selected pickets.

Another embodiment, although not the preferred one reduces the amount of DRAM memory available for each picket to allow inclusion of a section of fully associative memory of the type portrayed in Fig. 3. If say 512 bytes of fully associative memory were included, then every picket could contain a set of search indexes and in a single operation, 512 times 1024 pickets yields 512k compares per operation or 512 Gig-comparisons per second at one microsecond per operation. With extensibility, the concept can be advanced 40 into the multiple Tera-compare range. This embodiment empowers associative tasks that involve extensive searching for information with capabilities well in excess of computing capabilities today.

When this association operation using memory and byte wide coupled processing elements, as illustrated in Fig. 2, in addition to the applications of distinct algorithms or operations, artificial intelligence, 45 and parallel programming attempted in SIMD situations, there are many additional applications now available for the machine with the configuration of the chip we describe in a SIMD environment which may include:

- Simple parallelizable arithmetic tasks, including matrix multiply and other tasks which can be performed in specialized memory machines;
- 50 Image matching, and image processing tasks which can be performed in Von Neumann machines but which can be speeded up considerably with the applications adaptable to the extreme parallelism, for instance pattern matching of a three dimensional image;
- Data based query functions;
- Pattern matching in the Artificial Intelligence arena;
- 55 Network control in bridges to rapidly identify messages that go to a user on the other side of a bridge of a network;
- Gate level simulation;
- and checkers for VLSI ground rules violations.

Process tasks which take advantage of the bank of memories and associated processing elements will occur to application programmers as they invent to take advantage of the power of the new system architecture.

A process of keeping a description of a digital system can be enhanced by the use of the array to one gate or logic element per picket 100. In such a system, the process would begin by assigning each gate description as a list of signals that the gate accepts as inputs and naming the signal it generates. Require that each time a signal changes, its name is broadcast on bus 103 to all pickets and is compared in parallel with the names of expected input signals. If a match is found, record in the picket a new value of the signal in a dataflow register bit. When all signal changes have been recorded, cause all pickets to read out in parallel a control word which tells their data flow how to use the current set of inputs to compute the output. Cause these computations to be performed in parallel, with the results compared with the old value from the local gate. Record in a dataflow status bit all of those gates of the pickets whose outputs change. Cause an external controller to interrogate all the pickets and ask for the next gate that changed. Then broadcast the appropriate signal name and value from the picket to all other pickets, as originally stated, and repeat the cycle until no more signal changes occur or the process is stopped.

Another process would be a dictionary name search. Names are stored in picket memory 102 such that the first letter of all names can be compared with that of the desired broadcast name on broadcast data address bus 103. All pickets without a match are turned off with the control characteristic we provide. Then the second letter is compared and the compare and turnoff procedure is repeated for successive letters (characters) until no active picket units remain or the end of the word has been reached. At this point the remaining picket units are queried, and the index of the desired data is read out by the sequencer.

Fig. 4 illustrates a basic picket configuration of a plurality of parallel processors and memories, picket units, arranged in a row on a single silicon chip as part of a parallel array which may be configured as a SIMD subsystem, illustrating the control structure of such a system. Therein also is illustrated the control processor and the supervisory microprocessor. In Fig. 4, there is illustrated memory and parallel processing element logic on the same chip which within the Fig. 4 are shown in the section labeled Array of Pickets. Each memory is n bits wide, preferably a character wide, 8 (9) bits, as we have said, but conceptually also having a word width of multiple bytes wide memory. Thus, the memory portion of the parallel picket processor element will be preferably 8 (9) bits wide, or alternatively, 16, or 32 bits wide. With current CMOS foundry technology, we prefer to use an 8 bit or character wide associative memory (9 bit wide byte with self checking) with each picket processing element. The memories are directly associated one-on-one with a coupled processing element which includes an ALU, mask registers (A & Q used for masking), and a latch 104 (SP in FIG. 4), as well as status registers 107 and data flow registers A 105 and Q 106 (DF in FIG. 4) which are illustrated in greater detail in the Picket diagram of Fig. 2. The DRAM and logic of each picket processor do not have any burden of an interconnection network to contend with, for there is a direct one-on-one association between the multi-bit wide DRAM memory and its processing element on the chip itself.

It will be noted in Fig. 4 that the slide B register latch (SR) 104 is placed logically between the memory and the associated logic of the ALU of the processing element, and the latch becomes in essence a coupling port for each processing element along the picket array. Each picket chip comprises a plurality of parallel picket processing elements arranged in a line (which is illustrated as a straight bus) for communication with the picket control. A vector address bus is common to the memory, and a data vector address register controls what data passes across to each memory.

Fig. 4 also illustrates the interconnection between the main or microprocessor card MP which in our preferred embodiment is a 386 microprocessor configured as a PS/2 system with a subsystem controller through which global instructions pass to a canned routine processor CRP which we have provided which provides the instructions for an instruction sequencer 402 and execution control 403 which executes the specific microcode called for by the instruction sequencer. This instruction sequencer might be analogized to a controller in function. However, also within the canned routine processor CRP we have provided local registers 405 which together with the local reg's ALU (not shown) provide the basis for all addressing that is broadcast to all of the pickets within the picket array 406. In this way, address calculations are performed for all of the pickets in one ALU without using picket resources or conceivably without using picket execution cycles. This important addition adds control flexibility to the picket array, permitting dose, inhibit, and other control functions for specialized tasks to be performed and permitting pickets to be separated from any broadcast instruction or data function.

The instruction sequencer 402 with the loaded microcode 407 broadcasts to the array of pickets for execution under the SIMD instruction sequence determined by the main program microprocessor MP and by the canned routines of the canned routine processor CRP run time library 408 to enable SIMD processing of data contained in the array of pickets.

The instructions provided to the microprocessor MP via the subsystem interface are conceived to be high level process commands that might include Start Process, Write Obser. and Read Result which are passed to the microprocessor by the Subsystem controller of the microprocessor MP. The microprocessor can be considered as the main system or control processor in the subsystem arrangement illustrated in Figures 4, 5, 6 and 7. It will be understood that this unit could also be a stand alone unit with the addition of a peripheral input device (not shown) such as a keyboard and display unit. In this stand alone configuration, the system MP could be considered to be a commercial PS/2 to which cards which include the sequencer card (which configures the canned routine processor) and the processor array cards are inserted along the lines illustrated in Fig. 7. The routine library 411 can contain routine sequences for overall control of a process, such as CALL (.); Kalman, Convolve, and Nav. Update. The selection of these routines is via the user program, and thus the overall processing can be under the control of an external host or under control of the user program 412 located in the MP. A data buffer 413 is provided in MP memory for data transfers into and out of the parallel picket processor system. The instruction 402 sequencer is configured to execute the control stream from the MP as well as the canned routines that reside in the canned routine run time library memory 408. Some of these routines include CALL (.), Load Block, Sin, Cos, Find, Min, Range Comp. and Matrix Multiply from canned routines provided by the canned routine runtime library 408.

Within the CRP is also microcode 407 for the execution control of lower level functions like Load, read, Add, Multiply, and Match functions.

We prefer and provide an external FOR/NEXT control for each processing unit. We also provide a
20 deterministic floating point byte normalize implementation.

The use of a deterministic approach to macro development for the system, which we have provided, permits picket grouping and GROUP control. A local doze function is provided to accommodate individual picket processing variations.

Should the user program require execution by the array of processors, primitive commands, addresses and broadcast data is provided to the array of picket processors.

The particular function which each part of the system utilizes is determined by the task to be performed, and assigned during compilation of the user program.

The flexibility of the subsystem may be illustrated by a rather common problem. Take for example a matrix multiplication problem... $\{x\} * \{y\} = \{z\}$.

30 This would be described as the problem of:

$$\begin{array}{ccccccc}
 & & M & & C & & \\
 & x_1 & x_{R+1} & \dots & y_1 & y_{M+1} & \dots \\
 35 & \cdot & \cdot & & \cdot & R & \cdot & \cdot & X & M & \cdot & \cdot \\
 & \cdot & \cdot & & \cdot & \cdot & & & & & & \\
 & x_R & x_{2R} & \dots & x_{R \times M} & & & & & & & \\
 40 & & & & y_M & y_{2M} & \dots & y_{M \times C} & & & & \\
 & & & & & & & & & & & \\
 & & & & C & & & & & & & \\
 45 & & & & & & & & & & & \\
 & & & & z_1 & z_{R+1} & \dots & & & & & \\
 & & & & \cdot & \cdot & & & & & & \\
 50 & & & R & \cdot & \cdot & & & & & & \\
 & & & = & & & & & & & & \\
 & & & & \cdot & \cdot & & & & & & \\
 & & & & z_R & z_{2R} & \dots & z_{R+C} & & & &
 \end{array}$$

Which would be solved by the following statement adjacent which is illustrated the number of passes and number of clock cycles per pass is stated offered by way of example.

cycles/
passes pass 01 Call Matrix Mult Fx

```

1      c
5      (R,M,C,Xaddr, Yaddr, Zaddr)02  xSUB = ySUB = zSUB = 1
1      3 03  DO I = 1 to C              1      3 04  DO J = 1 to R
C      3 05      z = 0                  CxR      5/6* 06  DO K = 1 to
10     M      CxR      3 07  ***Assign to associative parallel processor***
08     Zz = Xx x Yy + Zz                CxRxM      204/345* 09
***Return result*** 10      xSUB = xSUB + R                CxRxM 2 11
ySUB = ySUB + 1                CxRxM      2 12  NEXT K
15     CxRxM      3 13  xSUB = xSUB - MxR + 1                CxR      2 14
ySUB = ySUB - M                CxR      2 15  zSUB = zSUB + 1
CxR      2 16  NEXT J                CxR      3 17  xSUB = 1
C      2 18  NEXT z                C      3 19  END Call
20     1      1

```

Note * Fixed Point (4 byte)/ Floating Point (1+4 Byte)

From a review of the above example, it will be seen that the task identified by the above statement 08 requires about 98% of the cycle time. Accordingly, it is assigned to the SIN organization of the parallel picket processor. The other processes take only 2% of the cycle times, and are maintained in the architecture within the microprocessor.

Accordingly, a review of this Example of matrix multiplication would be assigned for execution (and each statement upon compilation would cause execution in a specific system location) to either the MP, The CRP, the LR, or to the Picket array.

In the above example of matrix multiplication, the statement 01 would be assigned to the main processor MP, while statements 02, 05, 10, 11, 13, 14, 15, and 17 would be assigned to the local registers LR, while statements 03, 04, 06, 12, 16, 18, and 19 would be assigned for execution within the canned routine processor CRP, with the otherwise time consuming matrix processing thus assigned for execution under the single instruction to the array of pickets with assignment of the statement 08 to the array of pickets.

Fig. 5 represents a multiple parallel picket processor system 510 which incorporates a plurality of parallel picket processors. For applications such as multiple target tracing, sensor and data fusion, signal processing, artificial intelligence, satellite image processing, pattern/target recognition, Reed Solomon encode/decode operations, we have created a system which can be configured in a preferred embodiment as a SIMD system with 1024 parallel processors with two to 4 SEM E cards 511 (here represented as 4 cards per system) for each 1024 processors. The individual cards 512 are insertable in the rack mount system compartment 513 with wedgelock slides 514 and the cards are provided with insertion/extraction levers 516 so that when the cover 517 is closed it effectively encloses a mountable system in a rack with 32 or 64 M byte storage, with a performance capability of approximately 2 billion operations per second. The System is compact, and an array of a plurality of pickets is plugged into a backpanel board 518 which has logic thereupon which allows interconnection of a plurality of cards. The processor with 32 M byte storage is formed on 4 SEM E cards, and the system weighs only about 30 pounds. Power is provided by power supplies 519 as illustrated. Power requirements for the air cooled processor of such power is estimated to be only about 280 watts. Each SIMD system has two I/O ports 520 for channel adapter communication to associated mainframes or otherwise to the rest of the world. With the illustrated multiple parallel picket processor each consisting of 4 logic pages and using standard modular avionics packaging and bus structure for attachment to an external memory (e.g. PI, TM and IEEE 488 busses), the processor may attach via the I/O port to the memory bus of a mission processor and may be viewed as an extension of a mission processor memory space.

With the illustrated multiple parallel picket processor comprising 1024 parallel processing elements, each processor has 32kbytes of local memory and the associated path to the picket parallel processor is 8 bits or character wide (9bit) parallel.

Processors within each picket exchange data with other neighbor processors and between pages via a backplane interconnection network, preferably a crossbar, but alternatively a Slide Crossbar, a Shuffle Network, a Base 3 N-Cube or a Base 8 N-Cube.

Individual picket processors of the system are contained within a two card pack of the four cards, and the PS/2 microprocessor on one card, while the canned routine processor sequencer is contained on the other of the four cards which comprise the System schematically illustrated in Fig. S 6 and 7. Individual pickets 100, or cards of pickets 512 can be configured on the fly with the canned routine processor CRP to enter or can drop out of operations based on data conditions as controlled by the latch 104 architecture and the local registers 405 which are coupled to the execution control of the sequencer card CRP 703. Thus, picket processors can independently perform the alignment and normalization operations which are associated with floating point operations.

Processors are controlled in parallel by a common sequencer as described herein. The sequencer card 703 contains the controller CRP of the picket processor and may cause the picket process to execute a single thread of instructions, coded to execute on the array of SIMD processors in byte sequential fashion similar to classical bit serial processing as well. The controller has three layers. Micro-control for the pickets is microcoded similar to modern processors, and is transferred in parallel to all of the pickets. The Microcontrol and the Pickets are synchronized to the same clock system CLK, so that the functions controlled by the sequencer can be executed in the same clock time. Feeding commands to the micro-control sequencer is the function of the canned routine processor. This sequencer card 703 is a hardwired controller that during most functions executes loop control commands, and recursively starts new micro-control sequences. This controller with its canned routine library 408 and with its looping function keeps the pickets well fed, and not command bound. The canned routine processor controller CRP contains a large collection of macros that are called by the main system, which in the sub-system acts as a primary supervisory picket controller. This is the top control system of the picket array. It is a 386 microprocessors that manages the activity of the array of pickets. At a given moment all pickets of the array may execute the same instruction, although subsets of processors may react individually to the control flow.

There are several variations of the individual reactions, so that local autonomy, by virtue of the byte control function for each picket (dose, inhibit, etc.) there is local autonomy which programming may take advantage of and which may be made under control of the system during compilation of a program.

In addition, as has been described, there is local memory addressing autonomy. The SIMD controller sequencer supplies a common address for all of the pickets to use. Each picket can augment that address locally to enhance it's ability to do data dependent memory accesses.

In addition, a picket can participate or not participate in the array activities, depending on local conditions.

With this characteristic, it is now possible to introduce the concepts of groups to SIMD processing, by providing a means for each picket to assign itself to one or more of several groups, and processing can proceed based on these groupings, where a change of configuration can happen essentially on-the-fly. In one embodiment, only one group or combination of groups can be active at one time and each executes the same SIMD instruction stream. Some operations require only working with a subset or group of pickets. Programming can take advantage of this capability. Local participation autonomy is geared to work that way. Obviously, the more pickets that are computing the better.

One way to increase the number of pickets that are participating is to allow each picket to execute it's own instruction stream. This is essentially MIMD within SIMD. Now it is possible to essentially configure the same SIMD machine as a MIMD system or a machine of still a different configuration. This is because it is possible to program the picket to operate with it's own sequence of instructions.

Because each picket can be made to have its own sequence, decoding a very simple set of instructions at the picket level is possible and this allows more extensive local processing to take place. The areas where this function is most likely to find initial application is in complex decision making, however simple fixed point processing will be another area of interest to programmers.

A simple such program would load blocks of a picket program, say not to exceed 2K into picket memory 102, and these can be executed when the SIMD controller card 703 initiates local execution through the execution control beginning at a specified xyz address. This would be continued as the controller either counts off so many clocks, or tests for a task complete signal by monitoring the Status Funnel (SF) Registers illustrated in Fig. 4.

The status funnel (SF Fig. 4) makes use of the latch 104 for each picket. Each picket has a latch 104 which can be loaded to reflect the status condition of the picket. The SIMD controller can test the collective value in these latches (one per picket) by monitoring an array status line. This array status line is the logical combination of the values from each of the picket status latches.

In the following example, assume we want to adjust a value that are bigger than 250 to be in the range $500 < x \leq 250$. The following routine would use the status funnel to detect that the task was accomplished.

```

10      If VALUE < 500 then TURN YOUR PICKET OFF
      STAT < - PICKET OFF CONDITION
      IF STAT FUNNEL = OFF then finished - - - -
      VALUE < - VALUE - 250
15      Repeat

```

So, the multiple parallel picket processors configuration can be configured in various manners, and as a SIMD processor. Such a SIMD machine in the preferred embodiment is programmed to execute a single thread of instructions in the classical manner, and coded to execute on the array of SIMD processors in sequential fashion similar to classical processors, under general control of a SIMD controller or sequencer. At the application level this is accomplished by vector and vector like instructions, and vectors can be processed within processors and across processor. Vector instructions can be added with macroinstructions, typically with 6 to 10 such instructions.

With such a preferred embodiment, the system will appear schematically as illustrated in the functional block diagram of the parallel processor subsystem illustrated in Fig. 6. Through the I/O ports of the system as controlled by the host interface control 413 the subsystem sequencer functions similarly to a SIMD program with high function macros controlling the functions of the processing elements. Memory addressing enables an 8 bit, byte wide, data flow, and modulo8 arithmetic logic is used for functions (logical, add, multiply and divide). Provision is made for floating point format and for autonomous picket operation with individual sleep and dose mode and separate addressing.

The sub-system controller's arrangement is illustrated in Fig. 7. Each of the processor array cards 512 (shown as 4 in this subsystem illustration, but capable of being reduced to 2 SEM E cards) are coupled to a sequencer CRP 703, which is coupled to a subsystem controller 702 which in turn is ported to either the main memory system or to another subsystem in the configuration via a chip 705 interface to the associated microchannel bus 706. In the preferred embodiment the subsystem controller is a general purpose microprocessor unit of the IBM Corporations PS/2 (a trademark of IBM) and employs an Intel 386 processing chip and 4 Mbytes memory. The personal computer microprocessor MP 702 is coupled to the sequencer's card via a microchannel type bus 705 706 within the subsystem.

Obviously, many modifications and variations of this invention are possible in light of these teachings, and it is therefore understood that the appended claims permit the invention to be practiced other than as specifically described.

Especially, the following features may be combined alternatively or in common with one or more features of the appended claims:

- 1) the slide means enables transferring information in a single cycle to a non-neighbor picket unit position through a picket address location which would normally be able to receive the information were it not transparent to the message being sent until it arrives and is received at the nearest active neighbor which receives it;
- 2) the slide means enables a slide function by sending information to a non-neighbor position across "turned off" pickets enabling a first picket unit to transfer information to a remote picket when prior to a transfer cycle intervening pickets are made transparent by turning the intervening pickets units off, and then in the transfer cycle control means cause the first picket unit to send its information to its destination at a remote picket unit;
- 3) means are provided for transferring messages linearly across the lattice, across a two dimensional mesh, or in a three dimensional direction in an array;
- 4) the plurality of picket units have local autonomy, as well as being configurable as a coupled unit array of said picket units;

5) said system is configurable as a SIMD and MIMD system, and groups of said plurality of processing units are assigned to a programmable configuration, in which individual units of the system have a degree of local autonomy;

6) including a main processor system, said main processor system communicating with said external control sequencer across a bus, and said main system having means for causing global instructions to be executed by the parallel processing system;

7) said external control sequencer is coupled to a microcode memory, said microcode memory being programmable with canned routines for functions;

8) said sequencer is provided with high function macros for controlling the functions of the processing units which are coupled to said sequencer with a bus, and wherein memory addressing of the local memories of the system enables byte wide data flow and modulo 8 arithmetic logic to be used for logical, add, multiply and divide functions, and wherein provision is made for floating point operation within said parallel units, and for individual sleep and doze modes with separate addressing of individual picket processing units;

9) the operation of functions is programmatically allocable to be performed within said main processor system, said external control sequencer with canned routines, said local register means, or within said plurality of picket processing units, and wherein single instructions requiring extended processing of multiple data are assigned to said plurality of processing units which is configured for SIMD processing;

10) processing units of the system are coupled in an array of processing units, with processing units of the array having programmable local autonomy, and wherein a picket processing unit may operate with its own sequence of instructions, and can enter or drop out of operations connected with other processing units based on data conditions, and wherein processing units of the system can independently perform alignment and normalization operations which are associated with floating point operations;

11) said external control sequencer is configured to perform operations like CALL (), Load, Block, Sin, Cos, Find, Min, Range, and Matrix Multiply by canned routines provided by a canned routine runtime library;

12) canned routines are provided for execution control, Load, Read, Add, Multiply and Match functions;

13) a processing element exchanges data within an array of neighbor processing units of said array and between pages within said system via an interconnection network;

14) there is provided an external control processor for the picket array, and wherein micro-control code is transferred in parallel to all of the pickets in a group of said picket processors in said array, and wherein the control processor and the picket processing units are synchronized to the same clock system, such that the functions controlled by the external control processor can be executed in the same clock time;

15) above the external control processor for the picket array there is a top control system microprocessor coupled to the external control processor via a microchannel bus, which top microprocessor manages the activity of the array of pickets, and wherein the system is coupled so that at a given moment all picket processing units of the array may execute the same instruction, although subsets of processors may react individually to the control flow;

16) a plurality of picket parallel processing units are arranged along an address bus for communication with an external picket controller, and wherein there is a vector address common to the local memories of the picket system and wherein there is data vector address register means provided for control of what data passes to each local memory of the picket system. Communication with the picket control. A vector address bus is common to the memory, and a data vector address register controls what data passes across to each set associative memory;

17) there is provided a plurality of picket units arranged in an array and paths for data flow between the picket units having communication paths providing methods and means for bit parallel communication with all other pickets in the array while providing a coupling path to pass messages from one picket to any other;

18) messages can be passed from one to any other via an N-dimensional array implemented using Wire-or or Dot connected Pickets;

19) a small section of local memory be implemented such that every location in that section will simultaneously perform a match compare against a given pattern.

Claims

1. A parallel processing system comprising a plurality of picket units, each picket unit having a bit parallel processing element combined with a local memory coupled to the processing element for the parallel processing of information in all picket units in an associative way where each picket unit is adapted to perform one element of the association process.
2. A parallel processing system according to claim 1 wherein the memory of the plurality of picket units are arranged in an array and the array of pickets with their memory comprises a set associative memory.
3. A parallel processing system according to claim 1 or 2 further comprising a plurality of picket units formed on a silicon base chip and arranged in an array with neighbors for picket units to the side of adjacent picket units such that on a base chip there is formed a picket processing array with a plurality of memories, one for each of the processing elements capable of performing as a byte wide column content addressable memory arranged in a logical row or array with neighbor communication paths for passing data bidirectionally between the picket units and externally thereto.
4. A parallel processing system according to one of claims 1 to 3 wherein there is provided a picket processor array memory chip with a plurality of picket units and memory and having data flow paths between each picket's processing element and memory and between picket units.
5. A parallel processing system according to one of claims 1 to 4 where there is provided a plurality of picket units arranged in an array and paths for data flow between the picket units having one-on-one memory with each processing element of the array is across, left or right to or from an adjacent picket unit neighbor and with slide means providing for a slide communication with picket units farther away.
6. A parallel processing system according to one of claims 1 to 5 where each picket unit has a processor which has access to its own coupled local memory, and wherein character wide, or character multiples wide data and instructions flow between picket units in one clock cycle of the system.
7. A parallel processing system according to one of claims 1 to 6 wherein each picket chip has its own local memory providing character wide set associative storage in an array of picket units with at least 32 Kbytes storage provided for each local memory and there are sixteen picket units provided as nodes of a linear sub-array.
8. A parallel processing system according to one of claims 1 to 7 wherein there is provided a broadcast bus for communication between picket units which is multi-byte wide, so that all pickets can see the same data at the same time, and wherein there is provided means for passing picket control and address propagation transfers on this broadcast bus.
9. A parallel processing system according to one of claims 1 to 8 wherein local memory of a picket unit could be DRAM CMOS memory in a memory array and which supports row-column access by deleting the column demultiplexing on the back of the memory array, and which provides a row address that reads out a row of the memory array to cause data flows in parallel.
10. A parallel processing system according to one of claims 1 to 9 wherein the memory, in addition to data, contains "tri-bits" or "trit", so that there are three states recognized by the logic, signifying, either logic 1, logic 0, or don't care, and a trit is contained in successive storage locations in the storage array of the set associative memory.
11. A parallel processing system according to one of claims 1 to 10 wherein there is provided picket unit control means for providing a control function for individual operation by a picket unit.
12. A parallel processing system according to one of claims 1 to 11 wherein the local memory has a multibit binary reference storage address.

13. A parallel processing system according to one of claims 1 to 12 wherein is provided an external control store and control means for control functions which within an individual picket unit suspend operations in a picket unit which has a status output which meets a specific condition, which control functions provide a doze function, and an inhibit function, and an enable write to memory based on conditions in the picket unit and control functions which are provided to the picket unit after retrieval from said external control store.
14. A parallel processing system according to one of claims 1 to 13 wherein each local memory and processing element of a picket unit are provided with byte transfer means, and means for input of data and a mask for location of information in memory, and wherein means are provided for addressing memory in order to perform a search, and wherein there is means providing for input of data and a mask for a search in order to locate a word among N words in memory, and wherein matching locations raise a match line, and a separate operation reads or selects a first match, and wherein there are broadcast means for a broadcast operation in which a raised select line indicates participation and broadcast data is copied to all selected word locations.
15. A parallel processing system according to one of claims 1 to 14 wherein a processing unit comprises an ALU, mask registers, and a latch, as well as status registers (SR) and data flow registers (DF) which are coupled to a one-on-one local memory for each processing unit.
16. A parallel processing system according to one of claims 1 to 15 wherein local memory is a multi-bit wide DRAM and logic of each picket processing unit is formed on the same silicon chip substrate as the DRAM local memory, and wherein there is a direct one-on-one coupling between the local memory and its processing element, said local memory having cells which have a multi-bit address.
17. A parallel processing system according to one of claims 1 to 16, wherein each processing unit is provided with mask registers and a latch, said latch being placed to function as a coupling port for each processing unit along a communication line which is common to said plurality of processing units.
18. A parallel processing system according to one of claims 1 to 17 having an external control sequencer and local control register means for controlling the status of individual picket processing units of the system.
19. In a parallel processing system having an external controller and a plurality of processing units with logical gates and registers and memory contained therein, a process for keeping a description of a digital system comprising the steps of:
- assigning each gate description of the digital system as a list of signals that the described gate accepts as inputs and naming the signal it generates,
 - requiring that each time a signal changes, its name is broadcast to all processing units and is compared in parallel with the names of an expected input signal,
 - continuing to determine if a match is found, and recording record in the a processing unit a new value of the signal in a dataflow register, and
 - continuing until all signal changes have been recorded, and then causing all processing units to read out in parallel a control word which tells their data flow how to use the current set of inputs to compute the output,
 - causing these computations to be performed in parallel, with the results compared with the old value from the local gate, and
 - recording in a dataflow status register all of those gates of the the processing units whose outputs change, and
 - causing the external external controller to interrogate all the processing units and asking for the next gate that changed and

then broadcasting the appropriate signal name and value from the processing unit to all other processing units and repeating the cycle until no more signal changes occur or the process is stopped.

- 5 20. A process for using a parallel processing system having a plurality of processing units and local memories coupled for the transfer of data and controls there between comprising,

10 storing dictionary names in a processing unit memory such that the first letter of all names can be compared with that of a desired broadcast name broadcast to a plurality of processing units of the parallel processing system, and

15 wherein all processing units without a match are turned off with the control characteristic provided, and then comparing the second letter of the name and the compare and turnoff procedure is repeated for successive letters (characters) until no active processing units remain or the end of the word has been reached, and then providing a query for said processing units and then causing the index of the desired data to be read out of the processing units.

20

25

30

35

40

45

50

55




FIG.1
Prior Art




FIG. 2

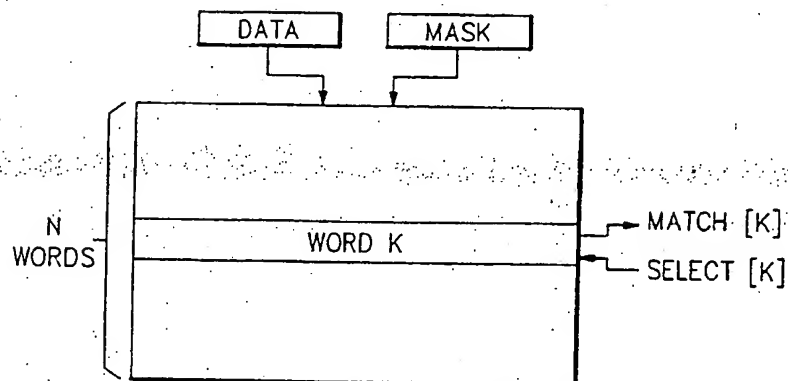


FIG. 3




FIG.4




FIG. 5

FIG. 6





FIG. 7





Europäisches Patentamt
European Patent Office
Office européen des brevets



Publication number: **0 485 690 A3**

EUROPEAN PATENT APPLICATION

Application number: 91109850.7

Int. Cl.⁵: **G06F 15/16, G06F 15/76,
G06F 15/80**

Date of filing: 15.06.91

Priority: 13.11.90 US 611594

Date of publication of application:
20.05.92 Bulletin 92/21

Designated Contracting States:
AT CH DE DK ES FR GB IT LI NL SE

Date of deferred publication of the search report:
21.09.94 Bulletin 94/38

Applicant: **International Business Machines
Corporation**
Old Orchard Road
Armonk, N.Y. 10504 (US)

Inventor: **Dieffenderfer, James Warren**

396 Front Street

Owego, New York 13827 (US)

Inventor: **Kogge, Peter Mivchael**

7 dorchester Drive

Endicott, New York 13760 (US)

Inventor: **Wilkinson, Paul Amba**

544 Forest Hill Rd.

Apalachin, New York 13732 (US)

Inventor: **Schoonover, Nicholas Jerome**

Box 18, Route 17C

Tioga Center, New York 13845 (US)

Representative: **Schäfer, Wolfgang, Dipl.-Ing.**
**IBM Deutschland Informationssysteme
GmbH**
Patentwesen und Urheberrecht
D-70548 Stuttgart (DE)

Parallel associative processor system.

Multiprocessor parallel computing systems and a byte serial SIMD processor parallel architecture is used for parallel array processing with a simplified architecture adaptable to chip implementation in an air cooled environment. The array provided is an N dimensional array of byte wide processing units each coupled with an adequate segment of byte wide memory and control logic. A partitionable section of the array containing several processing units are contained on a silicon chip arranged with "Picket"s, an element of the processing array preferably consisting of combined processing element with a local memory for processing bit parallel bytes of information in a clock cycle. A Picket Processor system (or Subsystem) comprises an array of pickets, a communication network, an I/O system, and a SIMD controller consisting of a microprocessor, a canned routine processor, and a microcontroller that runs the array. The Picket Architecture for SIMD includes set associative processing, parallel numeri-

cally intensive processing, with physical array processing similar to image processing. a military picket line analogy fits quite well. Pickets, having a bit parallel processing element, with local memory coupled to the processing element for the parallel processing of information in an associative way where each picket is adapted to perform one element of the associative process. We have provided a way for horizontal association with each picket. The memory of the picket units is arranged in an array. The array of pickets thus arranged comprises a set associative memory. The set associative parallel processing system on a single chip permits a smaller set of 'data' out of a larger set to be brought out of memory where an associative operation can be performed on it. This associative operation, typically an exact compare, is performed on the whole set of data in parallel, utilizing the Picket's memory and execution unit.

EP 0 485 690 A3

EP 0 485 690 A3




FIG. 1
Prior Art



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 91 10 9850

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.5)
A	JOURNAL OF VLSI SIGNAL PROCESSING, vol.1, no.1, August 1989, DORDRECHT NL pages 69 - 84 R. LEA 'ASP modules : cost-effective building-blocks for real-time DSP systems' * page 71, left column, line 9 - line 33 * * page 73, left column, line 25 - page 77, right column, line 15 *	1-18	G06F15/16 G06F15/76 G06F15/80
D,A	IEE PROCEEDINGS, vol.136, no.3, May 1989 pages 197 - 204 C. JESSHOPE ET AL 'Design of SIMD microprocessor array' * page 199, left column, line 7 - page 201, right column, line 48 *	1-18	
A	GB-A-2 223 867 (THE CITY UNIVERSITY) 18 April 1990 * abstract; claim 1; figures 3,5 *	1,5	
			TECHNICAL FIELDS SEARCHED (Int.Cl.5)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 1 March 1994	Examiner MICHEL T.G.R.
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application I : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1501 (3/82) (PM/COI)



European Patent
Office

CLAIMS INCURRING FEES

EP 91109850.7

The present European patent application comprised at the time of filing more than ten claims.



All claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for all claims.



Only part of the claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid.

namely claims:



No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirement of unity of invention and relates to several inventions or groups of inventions.

namely:

See Sheet 3.



All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.



Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid.

namely claims:



None of the further search fees has been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims.

namely claims:

1-18



European Patent
Office

LACK OF UNITY OF INVENTION

EP 91109850.7

The Search Division considers that the present European patent application does not comply with the requirement of unity of invention and relates to several inventions or groups of inventions, namely:

Claims 1-18 Byte social SIMD processor with associative memory.

Claims 19 Storage and updating of a digital system's description on a parallel processor.

Claims 20 Search of a name in a dictionary implemented on a parallel processor.